

October 1997

Lawrence
Livermore
National
Laboratory



Science & Technology

REVIEW



University of California
Lawrence Livermore National Laboratory
Science & Technology Review
P.O. Box 808, L-664
Livermore, California 94551

100 R&D AWARDS 1997



Nonprofit Org.
U. S. Postage
PAID
Albuquerque, NM
Permit No. 853

Printed on recycled paper.



About the Cover

The seven technologies and a sample R&D 100 Award plaque illustrate Livermore's excellent track record in *R&D Magazine's* annual competition for the year's 100 most technologically significant new products. The awards attest that the work done toward mission-related goals has broad application and is highly valued by outside industry. The winning technologies, which bring the Laboratory's cumulative total to 68, are highlighted in this issue.



What Do You Think?

We want to know what you think of our publication. Please use the enclosed survey form to give us your feedback.

Electronic Access

S&TR is available on the Internet at <http://www.llnl.gov/str>. As references become available on the Internet, they will be interactively linked to the footnote references at the end of each article. If you desire more detailed information about an article, click on any reference that is in color at the end of the article, and you will connect automatically with the reference.

About the Review

Lawrence Livermore National Laboratory is operated by the University of California for the Department of Energy. At Livermore, we focus science and technology on assuring our nation's security. We also apply that expertise to solve other important national problems in energy, bioscience, and the environment. *Science & Technology Review* is published ten times a year to communicate, to a broad audience, the Laboratory's scientific and technological accomplishments in fulfilling its primary missions. The publication's goal is to help readers understand these accomplishments and appreciate their value to the individual citizen, the nation, and the world.

Please address any correspondence (including name and address changes) to *S&TR*, Mail Stop L-664, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or telephone (510) 422-8961. Our electronic mail address is hunter6@llnl.gov.

S&TR Staff

SCIENTIFIC EDITOR

Ravi Upadhye

MANAGING EDITOR

Sam Hunter

PUBLICATION EDITOR

Sue Stull

WRITERS

Arnie Heller, Sam Hunter, Katie Walter, and Gloria Wilt

ART DIRECTOR AND DESIGNER

Ray Marazzi

INTERNET DESIGNER

Kitty Tinsley

COMPOSITOR

Louisa Cardoza

PROOFREADER

Al Miguel

S&TR is a Director's Office publication, produced by the Technical Information Department, under the direction of the Office of Policy, Planning, and Special Studies.

Printed in the United States of America

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-52000-97-10
Distribution Category UC-700
October 1997

Science & Technology REVIEW

October 1997

Lawrence
Livermore
National
Laboratory

- 2 **The Laboratory in the News**
- 3 **Commentary by Jeffrey Wadsworth
A Measure of Excellence**
- Feature**
- 4 **Livermore Science and Technology Garner Seven
1997 R&D 100 Awards**
- 6 **New Interferometer Measures to Atomic Dimensions**
- 8 **Compact, More Powerful Chips from Virtually Defect-Free,
Thin-Film System**
- 10 **A New Precision Cutting Tool: The Femtosecond Laser**
- 12 **MELD: A CAD Tool for Photonics Systems**
- 14 **The Tiltmeter: Tilting at Great Depths to Find Oil**
- 16 **Smaller Insulators Handle Higher Voltage**
- 18 **Computer Storage Management Software: The Next Generation**
- 20 **Patents**



USEC to Go Private

President Clinton has approved proceeding with the privatization of the United States Enrichment Corporation (USEC). USEC manages the Laboratory's Atomic Vapor Laser Isotope Separation (AVLIS) project, which is intended to produce enriched uranium for commercial nuclear reactor fuel. USEC was created in 1992 by Congress to privatize federal uranium enrichment activities, which convert natural uranium to enriched uranium for reactor fuel. With Clinton's approval, USEC will now be moved to the private sector.

"The vision . . . is to move uranium enrichment out of the government and into the private sector while realizing a substantial return for the U.S. taxpayer," said USEC Chairman William J. Rainer. "This latest action moves us to the final stages of realizing that vision." The action also furthers the nation's largest technology transfer effort.

USEC officials estimate it will take approximately six months to convert to private ownership. Added Victor Lopiano, director of the AVLIS program at Livermore, "This is an important first step toward AVLIS deployment."

LLNL Pantex Sign Weapons Stewardship Pact

The Laboratory and Pantex have formally joined forces to develop and implement new production technologies for caretaking and dismantlement of nuclear weapons, as well as demilitarizing excess nuclear weapons pits.

At a formal ceremony in July, William Weinreich, general manager of the Pantex Plant in Amarillo, Texas, and Livermore Director Bruce Tarter signed a memorandum of mutual intent for collaborative activities between the two Department of Energy facilities. The five-year agreement formalizes collaborative efforts long-ago established. Said Tarter, "Our role is to do the technology and work to transfer it to the production plants. It's important for all labs within the DOE complex to share their capabilities and their facilities."

The memorandum focuses on three areas: development of evaluation techniques and tools to assess the effects of aging on the nuclear stockpile and to assure the stockpile's reliability, technology to enhance the manufacturing of new high explosives that would replace the current (and aging)

stock of explosives should they be needed, and Laboratory-developed security systems that can be introduced at the Pantex site. The agreement also paves the way for an exchange of personnel and joint use of facilities.

Robotic Factory to Speed Up Human Genome Study

The Joint Genome Institute will be using robots to improve speed, accuracy, and economy of Lawrence Livermore, Lawrence Berkeley, and Los Alamos national laboratories in their work to decipher the human genetic code. The robots will read the order of DNA molecules and log the data into public databases in the process of unlocking the mysteries of human diseases. The robots will be installed in a 56,000-square-foot building in Walnut Creek, California, where the work will create up to 200 jobs.

Licenses for Microreaction-Chamber Technology

Lawrence Livermore National Laboratory has issued limited exclusive licenses to two diagnostic and research instrumentation firms for use of Livermore microreaction-chamber technology. The licenses granted to Cepheid of Santa Clara, California, and Soane BioSciences of Hayward, California, allow the companies to use the technology specifically for amplification and detection of nucleic acids and for ligand-binding assays.

The Livermore reaction chamber allows rapid changes and precise temperature control for very small reaction volumes (nanoliters) and is ideally suited for DNA amplification reactions. This capability is instrumental in clinical diagnostics and large DNA research studies such as the Human Genome Project.

The license to Cepheid permits the company to produce and sell products incorporating sleeve reaction-chamber designs that may include integrated optics for monitoring chemical reactions. In addition, Cepheid may produce and sell products that include electrophoretic devices coupled to an integrated optics reaction chamber.

The license to Soane BioSciences permits the company to produce and sell products that incorporate sleeve reaction-chamber designs without integrated optics, coupled to electrophoretic devices.



A Measure of Excellence

THERE are many markers of excellence for Laboratory technological achievement. One of the most telling, I believe, is the R&D 100 Awards that Livermore research teams receive each year. These awards showcase the very best of the world's new scientific and technical products and processes from corporations, universities, and government research centers.

The Laboratory's seven awards for 1997, featured in this issue, indicate the extremely high quality of creative scientific and engineering work that takes place in every Livermore program. The awards also demonstrate once again that the pursuit of our mission-related goals can advance U.S. economic competitiveness through the future development of important new commercial products for all citizens. Livermore's award-winning technologies, achieved by teams of researchers in highly specialized fields, represent significant contributions—in some cases fundamental breakthroughs—to society.

The diversity of our research programs contributes to the Laboratory's vitality. Our core strengths in nuclear science and technology, lasers and electro-optics, computer simulation of complex systems, advanced sensors and instrumentation, biotechnology, materials science, and advanced process and manufacturing technology spawn innovations in many scientific disciplines. These advances are prized by industry for their potential as the basis for improvements in existing products or for entirely new product categories.

Our broad expertise is reflected in the breadth of this year's R&D 100 awards and the fields they impact, including the semiconductor, opto-electronic, computer, oil, precision machining, and high-voltage energy industries. These industries have long been affiliated with Livermore researchers in one way or another. In fact, three of this year's awards are shared with our industrial partners. Working with

industrial partners helps us, on the one hand, to understand American industry's needs. One of our goals is to make the Laboratory even more accessible to industry partners and customers.

On the other hand, working with industry enables us to fulfill our national security mission and strengthen our technology base. The design and construction of the National Ignition Facility, for instance, are requiring extensive partnerships with industry to develop literally thousands of components, many of them unique.

By working with U.S. industry, Livermore "spin-offs" often directly benefit us through "spin-backs." For example, two Livermore R&D 100 awards this year are for the Absolute Interferometer and the Ultra Clean Ion Beam Sputter Deposition System. The first greatly improves the measurement accuracy of spherical and nonspherical optical surfaces, while the second significantly advances the state of the art in low-defect, thin-film deposition technology.

Together the two instruments appear to overcome serious hurdles that have blocked the advent of extreme ultraviolet lithography for making the next generation of powerful and compact computer chips. Once these chips are on the market, Livermore scientists and engineers will almost certainly use them to record, manage, and display their research data more effectively.

As a national security laboratory, Livermore was established to apply outstanding and innovative science and technology to pressing national problems. It is clear that our technological creativity is being recognized by the private as well as the public sector, as evidenced by our ability to win multiple R&D 100 awards this year and every year for the past 19 years.

■ Jeffrey Wadsworth is Deputy Director for Science and Technology.



Livermore Science and Technology Garner Seven 1997 R&D 100 Awards

THIS summer the Laboratory demonstrated once again its technical and scientific creativity by winning seven R&D 100 awards. The awards, given annually by *R&D Magazine* for the world's most technologically significant products and processes, are often called the "Oscars" of applied research.

Corporations, government labs, private research institutes, and universities throughout the world vie for the awards. Winners are chosen by the editors of the magazine and a large panel of experts in a variety of disciplines. The judges look for breakthrough products or processes that promise to improve people's lives through technological advances.

Past R&D 100 Award winners include products that have become fixtures of modern society—Polacolor film, electronic video recorder, antilock brakes, fax machine, halogen lamp, nicotine patch, color computer printer, and automated teller machine.

Since they began competing in 1978, Lawrence Livermore researchers have garnered 68 R&D 100 awards. Past Livermore winners include the precision engineering research lathe, ultralow-density silica aerogel, the hard-x-ray lens, the three-dimensional chemical x-ray microscope, the miniature mass spectrometer, and the electronic dipstick.

"The awards show that our mission-related work is clearly recognized for its strong potential to strengthen American industry and spawn innovative new products," noted Karena McKinley, director of Livermore's Industrial Partnerships and Commercialization office, which coordinated the entries.

The Laboratory's seven awards match its previous record for awards, set in 1987 and 1988. Four of the seven awards were won or shared by the Laboratory's Laser Programs, which has earned 34—or half—of the awards to date.

"This last year we challenged every directorate to make submissions," McKinley noted. "It is very exciting that this year's winners represent Engineering, Physics and Space Technology, Energy, Defense and Nuclear Technologies, Computation, and Laser Programs."

McKinley also expressed satisfaction that the past year has seen a stronger emphasis on collaboration. Some 41 Lawrence

Livermore scientists and engineers had their names on submissions, the largest number ever. Collaborators include three private companies (Veeco Instruments, Pinnacle Technologies, and IBM) and four DOE centers (Oak Ridge, Los Alamos, and Sandia national laboratories, and Allied Signal Federal Manufacturing and Technology Plant).

Livermore's award winning technologies are:

- **Absolute Interferometer** from Laser Programs. This instrument significantly improves the measurement accuracy of spherical and nonspherical optical surfaces. Today's commercial instruments are accurate to between a fiftieth and a twentieth the length of a visible light wave. However, this instrument features an accuracy of nearly a thousandth of a visible light wave (about the width of one atom). The instrument uses diffraction to generate and then compare two perfectly spherical light waves, one as a reference and one to illuminate the part being measured. This new instrument will expand the frontiers of the semiconductor, optical, and metrology industries. It will also help make possible the high-resolution optical systems that are required for extreme ultraviolet lithography.

- **Ultra Clean Ion Beam Sputter Deposition System** by teams from Veeco Instruments and Laser Programs. This high-precision instrument significantly advances the state of the art in low-defect, thin-film deposition technology. The system produces precise, uniform, thin films on substrates, reducing defects by a factor of 100,000. The tool is applicable to a large class of thin-film coatings used in current semiconductor production and very low defect-density films needed for ultrahigh-density multilayered magnetoresistive heads for the magnetic recording industry. The achievement is also a significant step toward realizing the potential of extreme ultraviolet lithography for manufacturing the next generation of computer chips.

- **Femtosecond Laser Materials Processing** from Laser Programs and Defense and Nuclear Technologies. This new tool features full computer control of a laser to accurately machine all kinds of materials (i.e., steel, ceramic, heart and

dental tissue) with minimum waste. The machine uses extremely short laser pulses so that material is removed atom by atom. As a result, material even within 0.1 micrometers of the machined surface is not damaged. Because the process does not require melting or boiling, the laser can cut and drill objects that cannot be machined by conventional techniques. In addition, the technology permits the machining of slots and holes down to about 1 micrometer in size.

- **Multiscale ElectroDynamics (MELD)** from Engineering and Physics and Space Technology researchers. This simulation software has the potential to revolutionize the design process for opto-electronics by creating a "virtual optical bench." MELD models very different elements, such as semiconductor waveguides, fibers, and lenses and provides a seamless interface between them. The software enables designers to quickly and accurately explore new designs and packaging approaches, eliminating trial-and-error methods, thereby drastically reducing development cycle time and nonrecurring engineering costs. By reducing fabrication cycles, optimization time, and cost, the software offers the potential to increase the U.S. share in the worldwide opto-electronic component market.

- **Oil Field Tiltmeter**, a collaboration between Livermore Energy and Environmental Programs scientists and Pinnacle Technologies Inc. of San Francisco. This tool detects very slight changes in rock cracks down to 10,000 feet in oil wells. An array of these instruments is used to monitor oil-well hydrofracturing—a technique of cracking deeply buried rock to provide channels through which oil can flow. The information produced from the instruments is used to generate a map of the surface deformation around the oil well for choosing optimal sites for neighboring oil wells. Previous technology could monitor hydrofractures only 6,000 feet deep, a considerable limitation because about 80% of hydrofracturing is done deeper than that.

- **Ultra High Gradient Insulator** from Livermore's Defense and Nuclear Technologies and Laser Programs and the AlliedSignal Federal Manufacturing and Technology

Plant. This breakthrough insulator technology improves the ability of insulators to resist breakdowns by up to a factor of 4. It is constructed of ultrafine, alternating layers of insulating and conducting materials. The insulator's marked increase in performance should revolutionize linear accelerators and reduce the size and cost of high-voltage applications, including x-ray machines, neutron sources, power transmission equipment, and plasma radiation sources. The insulator will also permit smaller and entirely new types of high-voltage equipment that were not previously possible.

- **High-Performance Storage System**, a collaboration of Livermore, Los Alamos, Oak Ridge, and Sandia national laboratories and IBM Global Government Industry. This new system, based upon innovative scalable architectural concepts, increases both the performance and capacity of storage systems for large-scale computation by a factor of 100 or more. The system will help advance diverse applications such as modeling and simulation on supercomputers, storage and real-time delivery of motion pictures, medical image processing, oil exploration, high-energy physics, satellite data gathering, and the finance and insurance industries. In particular, the system will meet the needs of the DOE's Accelerated Strategic Computing Initiative, a major element of the Stockpile Stewardship and Management Program.

These seven R&D 100 Award-winning technologies and their creators are featured in more detail in the articles that follow.



For further information contact
Karena McKinley (510) 423-9353
(mckinley3@llnl.gov).

New Interferometer Measures to Atomic Dimensions

HOW do you accurately measure the surface of a mirror to check for high or low spots that are no larger than a few atoms? Until recently, you couldn't. But that has changed, thanks to a team at Lawrence Livermore. Their new Absolute Interferometer can measure large surfaces to find uneven spots less than 1 nanometer (billionth of a meter) higher or lower than the rest of the surface.

Optical interferometers are instruments that can make very precise measurements of objects using the interference pattern of two waves of light. One wave interacts with the object being measured, and the other does not; their interference when they encounter one another allows measurements to within one-thousandth of the wavelength being used. Very small distances and thicknesses can be measured, including extremely small surface irregularities in optical devices such as mirrors. In astronomy, interferometers are used to measure the distances between stars and the diameters of stars.

For your bathroom mirror, such perfection is hardly necessary. But for high-end optical applications, accuracy is essential. The semiconductor industry will find the new interferometer indispensable as the demand for ever more powerful microchips necessitates a change in chip printing methods.

As the tiny circuits printed on microchips are made smaller, more circuits and hence more information can be included. Today, the binary-circuit patterns are projected onto a resist-coated silicon wafer. The size of the features on the chip is limited by the shortest wavelength that the lenses in the projector will transmit. When the wavelength gets down to about 180 nanometers, no lens can transmit it.

To make the chips' features smaller, mirrors can be used to reflect rather than transmit the light, allowing the use of light with wavelengths as short as 13 nanometers. This new process is known as extreme ultraviolet (EUV) lithography because the light used is in the far edge of the ultraviolet range of the spectrum. With it, microprocessor features can be made as small as 0.1 nanometers, which is about 1,000 times smaller than the width of a human hair. With current lithographic methods, the smallest achievable feature size is

0.18 nanometers. Today, the smallest features in production are 0.35 nanometers.

To reflect such short wavelengths, mirrors must have high and low spots (known as surface figure errors) less than approximately 0.25 nanometers. Fabricating such a mirror requires surface measuring systems that border on perfection. This is where the Absolute Interferometer comes in. The brainchild of physicists Gary Sommargren, Donald W. Phillion, and Eugene Campbell and designer Franklyn Snell, the new interferometer represents a 100-fold improvement in accuracy for measuring surface shapes of optical components and removes one of the blocks to furthering the development of EUV lithography.

Adding Diffraction

Like all interferometers, this one uses the interference pattern of two waves of light to measure objects or phenomena. These light waves are usually imperfect because they are dependent on the condition of the surface or lens from which they emanated, and this imperfection introduces error into the measurements. To correct this problem and produce a nearly perfect spherical wavefront, Livermore's new interferometer incorporates diffraction, which is the breaking up of light as it passes around an object or through a hole. The core of an optical fiber acts as an aperture through which the light beam passes. There are other optical interferometers that incorporate diffraction, but this one is different in that the two wavefronts are generated independently. Their relative amplitude and phase can be controlled, yielding the contrast adjustment and phase-shifting capability necessary for the highest possible accuracy.

A frequency-doubled, Nd:YAG (neodymium-doped yttrium-aluminum-garnet) laser operating at a 532-nanometer wavelength launches light into two single-mode optical fibers. As shown in the figure (p. 7), the light diffracts on exit from the fibers, forming spherical wavefronts. The "measurement" wavefront passes through the optical system being tested, which induces aberrations in the wavefront and causes it to focus on the endface of the other fiber. Here the



Shown with the Absolute Interferometer are (back, left to right) Franklyn Snell, Donald Phillion, (front) Gary Sommargren, and Eugene Campbell.

wavefront reflects off a semitransparent metallic film on the fiber endface and interferes with the "reference" wavefront to generate an interference pattern. The pattern is recorded by a charge-coupled device camera. Associated software describes the magnitude and spatial distribution of errors so that correction strategies can be devised.

The design is simple, containing only the optic being tested and the two optical fibers that generate the two wavefronts. This design makes the interferometer versatile for measuring optical components and systems and allows it to measure in a single-pass transmission, unlike conventional interferometers. The critical component is the fiber endface coated with a semitransparent film. It must have a flatness comparable to the desired accuracy of the measurement, but only over a very small area around the fiber core. By embedding the fiber in a glass substrate and superpolishing the entire assembly, achieving the surface finish is easy. To ensure stability and ease of mounting, the fiber remains embedded in the substrate during use. Typical fiber cores have a diameter of 3 micrometers. The critical, carefully polished region has a diameter of about 1 millimeter.

The system's design assures its accuracy in at least two ways. The fibers act as spatial filters, correcting for any lack of quality in the wavefronts as they pass through the fibers. And before the two wavefronts interfere, they encounter no other optical components that can degrade accuracy, except for the one fiber endface.

Improved Optics

The Absolute Interferometer opens the door to fabrication of optical components for EUV lithography with surface

errors within approximately 0.25 nanometers. Fabrication of these mirrors requires real-time, surface-measurement feedback during the polishing process at a level of accuracy that conventional interferometers cannot provide.

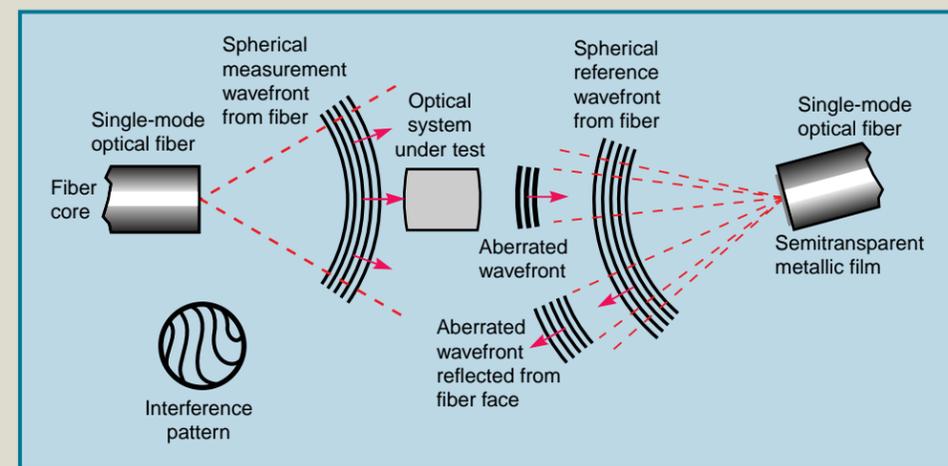
Livermore's new interferometer can be used to test any optical component or system where extreme accuracy is required. Perhaps more important, it can also be used to create higher-quality reference surfaces for use in conventional interferometers. These reference surfaces would advance the accuracy of the large base of interferometers already installed in the optical industry. Small companies would then be able to increase their capabilities without having to purchase new interferometers.

The usefulness of any optical system is limited by its inability to produce perfect images. As we move to extreme ultraviolet and soft-x-ray wavelengths for microscopic medical imaging or deep-space astronomical imaging, advances in optical fabrication and testing must keep pace. The Absolute Interferometer brings new accuracy to these applications.

—Katie Walter

Key Words: extreme ultraviolet (EUV) lithography, fiber optics, integrated circuit manufacturing, interferometry, optics, R&D 100 Award.

For further information contact Gary Sommargren (510) 423-8599 (gesommargren@llnl.gov).



This schematic (not to scale) shows how the Absolute Interferometer operates using single-mode optical fibers. In this example, the interferometer is testing an imaging lens, but the setup is similar for testing imaging mirrors.

Compact, More Powerful Chips from Virtually Defect-Free, Thin-Film System

THE prospect of computer chips ten times more powerful yet one-tenth the size of today's has taken a large step forward with the development of the Ultra Clean Ion Beam Sputter Deposition System, the result of a collaboration between Livermore's Laser Programs and Veeco Instruments Inc. of Plainview, New York. The collaborating team received a 1997 R&D 100 Award for their successful development.

The new machine, called the LDD-IBSD 350, deposits extremely thin single and multilayer film coatings with angstrom accuracy onto substrates of silicon and other materials. These films are used in creating advanced computer chips, hard disk drives, and the master patterns for extreme ultraviolet (EUV) lithography. The machine significantly advances the state of the art by drastically lowering the number of thin-film defects it deposits by a factor of 100,000 over existing equipment. Precise application of thin films is critical to the \$120-billion semiconductor manufacturing industry and the \$100-billion magnetic recording industry.

Compact but More Powerful

By vastly reducing the number of defects in thin films, the machine also eliminates one of the most significant obstacles to the realization of EUVL for the next generation of computer chips. EUVL will allow chip makers to work with wavelengths much shorter than today's deep ultraviolet light-based technology, reducing line widths (and feature sizes) to below a tenth of a micrometer (a millionth of a meter).

"The microchip industry now uses photolithography that employs wavelengths of light in the deep UV region [250 nanometers, or billionths of a meter]," explained Richard Levesque, Laser Programs engineer and award recipient. "To make chips more powerful and compact, we need to decrease the wavelength by using the extreme ultraviolet region of the electromagnetic spectrum [13.4 nanometers]."

Because conventional lenses and accompanying optical systems do not transmit EUV light, semiconductor

researchers have turned to reflective multilayer-coated optical systems that are designed to operate at a new standard of 13 nanometers.

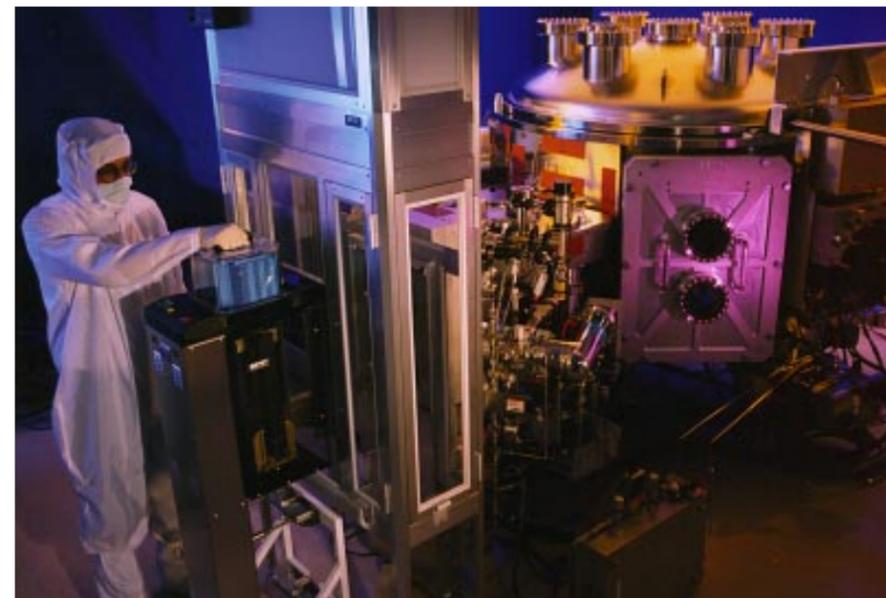
By using EUVL, manufacturers should be able to produce considerably more powerful chips yet shrink them to one-tenth their current size. Indeed, it is expected that the forthcoming technology will produce single chips containing 100 million transistors, truly the long-awaited "computer system on a chip."

Conquering Defects

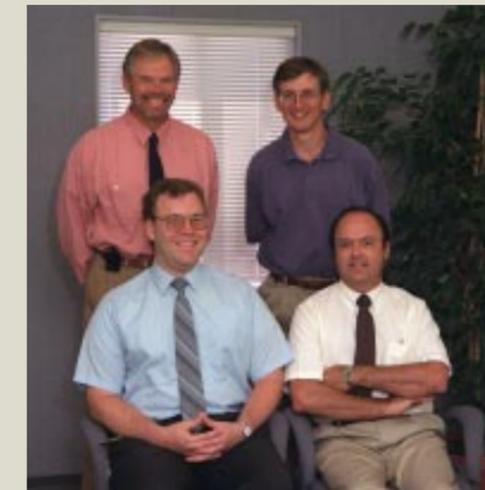
A critical problem remained, however, with the thin-film coatings used in the reflective mask (the master pattern used to "print" the semiconductor circuits onto silicon wafers or chips). Livermore team member and Advanced Microtechnology Program leader Don Kania calls masks the "Achilles' heel of all advanced lithography systems." In the case of EUVL, if the reflective mask coatings have even extremely small defects, they would be replicated, or printed, in the lithography process onto the computer chips being manufactured, thereby destroying the chips' complex circuitry.

The conventional approach for making masks, known as magnetron sputtering, produces about 10,000 defects in a square centimeter, far too many for successful EUVL. What's more, there are no viable techniques for repair or replacement of defective regions of multilayer coatings. As a result, development of EUVL technology has been contingent in part upon a "defect-free" multilayer deposition process. The development of the IBSD-350 process for film coating, with its remarkably low defect density, effectively solves this long-standing problem.

For EUVL applications, the machine produces precise, uniform, high-reflective masks with 81 alternating layers of molybdenum and silicon, each 3 to 4 nanometers thick, on 150-millimeter-diameter silicon wafers. Having excellent control of the thickness of the 81 layers in the multilayer is paramount to the optical performance of the mask. In depositing the layers, the machine directs a beam of ions at



Above, the Ultra Clean Ion Beam Sputter Deposition System coats silicon chips with ultralow-defect films at angstrom accuracies, allowing considerably more powerful chips at one-tenth current sizes. At right, Livermore team members include (left to right, standing) Steve Vernon, Don Kania, (sitting) Patrick Kearney, and Richard Levesque.



molybdenum or silicon targets. The ions physically collide with the target, forming a vapor, which is then precisely deposited on the substrate with a defect density of less than 0.1 per square centimeter.

The machine's defect density represents a 100,000-fold improvement over typical defect levels for multilayers produced in conventional physical deposition processes. It also represents a sixfold improvement over the goal of the 1998 Semiconductor Industry Association (SIA) for particle contamination of bare silicon wafers and meets the SIA goal for process-induced contamination for the year 2004.

The development of the machine culminated 16 months of high-risk, high-payoff effort that combined the ion-sputtering technology expertise of Veeco with Livermore expertise in high-vacuum technology, opto-electronics, micro-electronics, and x-ray diagnostics built up over the past two decades from research in defense and laser fusion.

In particular, Laser Programs' Advanced Microtechnology Program was a natural partner for Veeco because of the members' strong expertise in processes for fabricating the microstructures used as targets and diagnostic equipment in Livermore inertial confinement fusion experiments. Over the past few years, the program has been working on the challenges posed by EUVL.

The Livermore team consisted of physicist Steve Vernon, mechanical engineer Levesque, and materials scientists Patrick Kearney and Kania. According to Levesque, "Our counterparts at Veeco had a real willingness to work with us, share their expertise, and also were receptive to Livermore ideas."

The Veeco-Livermore development team achieved the new process by integrating a breakthrough film-deposition module

to a state-of-the-art wafer handling system, advanced vacuum technology, automated substrate handling, and mechanical systems. The machine permits flexibility in the type and number of thin-film-deposition processes that can be implemented by the end user.

Coatings for the Future

While the new technology helps to open the door to EUVL, its flexibility, uniformity of film deposition, and stable deposition rate make the machine a strong candidate for conventional processes used in current semiconductor production of a broad class of thin-film coatings where low defect density is a strong concern. Such coatings can be virtually any material or combination of materials including metals, semiconductors, and insulators. Near-term applications also include the fabrication of very low defect-density films for ultrahigh-density, multilayered magnetoresistive heads for the magnetic recording industry.

—Arnie Heller

Key Words: extreme ultraviolet (EUV) lithography, photolithography, R&D 100 Award, thin-film deposition, Ultra Clean Ion Beam Sputter Deposition System.

For further information contact Stephen Vernon (510) 423-7826 (vernon1@llnl.gov).

A New Precision Cutting Tool: The Femtosecond Laser

CHIRPED-pulse amplification strikes again. Using it in a high-peak-power mode, Laboratory scientists produced first the 100-terawatt laser and then the petawatt laser, opening up new opportunities for applying laser-matter interactions. Now a Livermore team has won an R&D 100 Award for applying chirped-pulse amplification in a high-average-power mode for cutting and machining materials. The system was developed for disassembling nuclear weapons components, but it has many other uses as well.

The team, led by Brent Stuart, illustrates Livermore's collaborative nature by combining research and development expertise from Laser Programs and Defense and Nuclear Technologies Directorates.

From Demilitarization to Dentistry

By ionizing the material being cut—removing it atom by atom—the cutting technique allows precise machining of everything from steel to tooth enamel to very soft materials like heart tissue. Each pulse of this machining system is extremely short, lasting just 50 to 1,000 femtoseconds (or quadrillionths of a second). These ultrashort pulses are too brief to transfer heat or shock to the material being cut, which means that cutting, drilling, and machining occur with virtually no damage to surrounding material. Furthermore, this revolutionary laser can cut with extreme precision, making hairline cuts in thick materials along a computer-generated path.

In dentistry applications, the thermal nature of the conventional laser ablation process can heat and crack a tooth and produce a random-shaped hole within a large area of collateral damage. In contrast, at the same ablation rate, Livermore's new laser precisely removes the material and leaves the surrounding areas in their original state (see a and b of the figure, p. 11).

The ultrashort-pulse laser represents a major advancement in cutting technology. Conventional lasers, diamond saws, and water jets are used commercially for a variety of cutting and machining applications. But each one has drawbacks. None of them can achieve the precision of the femtosecond laser machine tool (0.1 millimeters), and most of them damage surrounding material to varying degrees. Because of these shortcomings, no commercial cutting system can be

used on the range of materials or applications of Livermore's new tool.

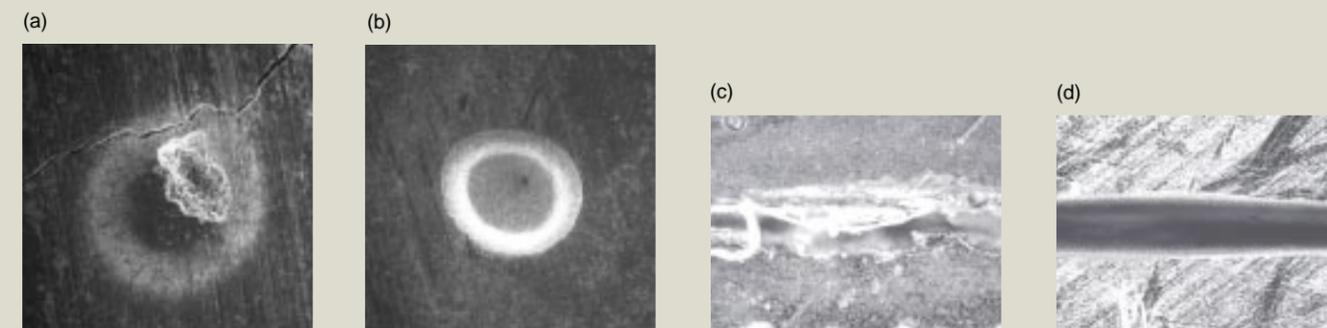
Industrial lasers, which melt and boil material to remove it, are often used in precision cutting. The heat and shock cause considerable damage to the area surrounding the cut that can range from changes in the grain structure to cracking. The damage may extend from a few micrometers to several millimeters from the cut, depending on the properties of the material, the laser pulse duration, and whether a cooling method is used. Very tiny structures only a few tens of micrometers in size, such as biological tissue or semiconductor devices, are extremely fragile. Even the slightest thermal stress or shock creates intolerable collateral damage.

These conventional cutting methods also leave slag around the cut. When material is vaporized, some of it is deposited on the walls or upper surface of the cut. This residue reduces the quality of the cut and the efficiency of the cutting system.

With each short pulse of the Laboratory's new laser cutter, material is heated to temperatures far beyond the boiling point, producing an ionized plasma, while leaving surrounding material cool. The pulse deposits its energy so quickly that it does not interact at all with the plume of vaporized material, which would distort and bend the incoming beam and produce a rough-edged cut. The plasma plume leaves the surface very rapidly, ensuring that it is well beyond the cut edges before the arrival of the next laser pulse. And because only a very thin layer of material is removed during each pulse of the laser, the cut surface is very smooth and does not require subsequent cleanup (see c and d of the figure, p. 11).

Removal of minimal amounts of material makes this new cutting system useful for processing extremely valuable or hazardous materials. If the cutting is done in a vacuum, better than 95% of the removed material can be recovered.

Another Livermore team is building a high-powered femtosecond machining system for the Department of Energy's Y-12 Plant at Oak Ridge, Tennessee, one of this country's primary manufacturers of nuclear weapon components. A second unit at Livermore will be used as engineering support to the Y-12 unit. The high precision of this cutter will maximize the plant's ability to reuse high-value components and minimize the amount of waste generated during the cutting process.



For cutting teeth, (a) a conventional laser cutter causes heating and cracking a result of large thermal stresses, while (b) the femtosecond cutter produces a clean hole with no collateral damage. In stainless steel, (c) a conventional infrared laser (wavelength of 1,053 nanometers) operating at a pulse duration of more than 1 nanosecond produced a jagged cut and much slag, but (d) Livermore's new cutter, with a pulse duration of 350 femtoseconds and the same wavelength, produced a clean cut with no slag.

Livermore is studying the use of the Femtosecond Laser to machine high explosives for experiments at its High Explosives Applications Facility. Because so little energy or mechanical shock is transferred to the part being machined, the team has demonstrated that materials such as high explosives or parts containing high explosives can be cut without danger of detonation. The team is also working on the design of a system for demilitarizing chemical weapons.

Other potential applications abound. Using the laser as a surgical tool for soft tissue has already been discussed in *Science & Technology Review* (October 1995, pp. 28–31). A semiconductor device producer is exploring the use of the unit for cutting high-value semiconductor wafers. Other major U.S. manufacturers are looking into incorporating femtosecond machining systems into their production lines. In manufacturing, new materials are constantly appearing, and the features on all kinds of devices are becoming smaller and smaller. The femtosecond machining system may be the most effective way to respond to both challenges with its high precision on all materials regardless of composition.

—Katie Walter

Key Words: chirped-pulse amplification, demilitarization, femtosecond laser, laser surgery, R&D 100 Award.

For further information on the technology and applications contact Brent Stuart (510) 424-5782 (stuart3@llnl.gov) or Michael Perry (510) 423-4915 (perry10@llnl.gov).



The team that developed the Femtosecond Laser includes (front) Paul Armstrong; (middle, left to right) Alexander Rubenchik, Hoang Nguyen, Steve Herman, and Brent Stuart; (back) Michael Feit, Booth Myers, Michael Perry, Joseph Sefcik, and Howard Powell (Paul Banks is not pictured).

MELD: A CAD Tool for Photonics Systems

COMPUTER-aided design, or CAD, has been used by architects and engineers for years to design everything from spas to the Space Shuttle. A CAD program called SPICE, used to design integrated circuits, helped to fuel the ongoing revolution in integrated circuit technology that has formed the backbone of the information age. But the information age is changing. Now, many micro-electronic and micro-optic systems combine to form photonic devices, which manipulate light for control, communication, sensing, and information displays. Photonics systems have been found in Lawrence Livermore projects since the late 1980s. Future applications include a nonlinear optical waveguide that is being designed for use in the National Ignition Facility.

Photonic circuits using integrated and micro-optic devices are expected to form the basis of all future high-speed and wide-bandwidth communication systems, computers, and signal- and image-processing hardware. Design of photonics devices, while challenging, is very important because the devices are expensive and difficult to fabricate. Accurate computational simulation for both design and packaging (alignment of the very tiny parts that constitute photonic devices) will reduce costs considerably and remove a major barrier to widespread use of photonic systems.

A team at Lawrence Livermore has created a computer code to allow designers to quickly and accurately explore new photonics design and packaging approaches. The code eliminates tedious trial-and-error fabrication and reduces development cycle time and engineering costs by as much as 80%.

This new software can handle the design of photonics systems with widely disparate component scales, hence the name Multiscale ElectroDynamics, or MELD. No other photonics design software has this capability. Designed by a team that includes Richard Ratowsky, Jeffrey Kallman, Robert Deri, and Michael Pocha, MELD has the potential to revolutionize the design process for photonic devices and packages.

MELD allows the user to construct a "virtual optical bench," shown in the monitor below, where photonics modules are placed. Each module treats a single photonics element and uses its own optical simulation algorithm. The code then provides for seamless interfaces, or links, between these algorithms, thus allowing simulation of the entire system. The links exchange information between modules, including accumulated transmitted and reflected optical fields. This system can be used for optically large structures such as waveguides, for optically small structures such as abrupt discontinuities in materials, or for optically mixed structures such as antireflection-coated lenses.

MELD is written in object-oriented C++, a programming language that allows easy integration of new modules.

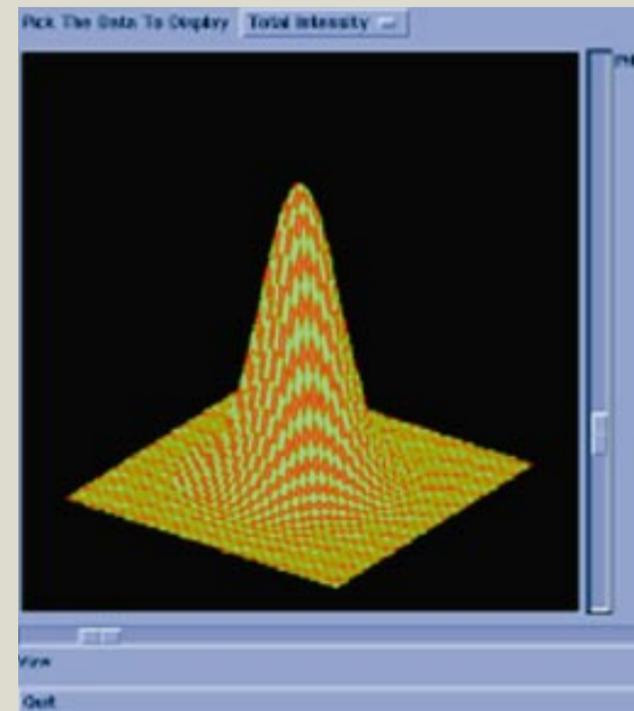
MELD's capabilities are valuable only if the resulting system is also accurate.

MELD has already proved its accuracy at Hewlett-Packard

Shown with the MELD "virtual optical bench" is the team that created it: (left to right) Richard Ratowsky, Michael Pocha, Jeffrey Kallman, and Robert Deri.



Company Laboratories in Palo Alto, California, where it was validated for the design of ball lenses, devices that match the mode of a laser diode to a single-mode optical fiber. The HP Labs' microphotonics effort is aimed at developing packaging techniques for low-cost, fiber-optic components, and their researchers had found that traditional methods of calculating the coupling efficiency for a laser into a fiber through a ball lens were not accurate enough. The spherical aberration of the ball lens strongly influences the apparent focal length of the lens and the achievable coupling efficiency. It is also important for optical subassemblies to minimize the coupling of reflections back to the laser. MELD addresses both issues



The designs of fiber-optic communication systems demand crucial calculations such as this one. Here, MELD displays the intensity profile of light propagated from a laser diode through a ball lens.

effectively. In fact, one of the HP Labs engineers has said that now when they find any discrepancy between the experimental results and the results indicated by MELD, they go back to check their experimental setup.

MELD has competitors already on the market, but none offers MELD's range of capabilities. Most can simulate optically large waveguiding structures, accurately treat aperture effects, and handle optically large components by ray tracing. MELD is unique in its ability to handle both large and small structures accurately; in particular, MELD invokes a true Maxwell equation solver when necessary. Maxwell's equations are a set of classical equations that govern the behavior of electromagnetic waves. When calculating radiation from small structures, the usual beam propagation techniques will not work, and an accurate calculation can be made only with a Maxwell solver. Solving Maxwell's equations is also needed for calculating reflections, a critical design factor for photonic systems. For example, reflections from ball lenses are often extremely small, but they can destabilize active devices such as laser diodes. MELD allows for the global calculation of reflections from a variety of elements.

The photonics component market today—in information and communications equipment, consumer electronics, and national security applications—amounts to about \$15 billion annually worldwide. By sharply reducing development cycle time and nonrecurring engineering costs, MELD offers the potential to increase the U.S. share of this growing market. For uses related to national security, flexible manufacturing through accurate design is particularly crucial, given the small production runs in today's downsized defense industry. Saving time and money, regardless of the situation, is always advantageous.

—Katie Walter

Key Words: computer-aided design, fiber-optic communications, flexible manufacturing, photonics, R&D 100 Award.

For further information contact Richard Ratowsky (510) 423-3907 (ratowsky1@llnl.gov).

The Tiltmeter: Tilting at Great Depths to Find Oil

TIME was when a wildcatter would cross his fingers, sink a well, and hope for a gusher. Now, with oil getting harder to find and explorers having to drill to 10,000 feet (3,000 meters), the cost of a well can run in the multimillion-dollar range. So when an oil exploration company is planning its next well, its management wants to rely on something a little more substantial than hope.

That's why the R&D 100 Award-winning Oil Field Tiltmeter has such a great future, enabling explorers to determine where to put their next multimillion-dollar well to get maximum production.

More than half of American oil-producing wells require help, usually in the form of cracking the underground rock to provide channels through which the oil can flow. This hydrofracturing is done by pumping a mixture of water, polymers, and sand down a well under high pressure, which causes the surrounding rock to crack and move slightly. In the past, an array of ultrasensitive tiltmeters was placed near the surface to detect the slight tilting, which reveals the primary direction of the cracking several thousand feet below. This information helps drillers decide where to sink additional wells.

But these surface tiltmeters were not without problems. They could measure hydrofracturing only to a depth of 6,000 feet (1,800 meters), which excluded about 80% of the oil wells in the nation. They also were vulnerable to "noise" at the surface—wind, thermal expansion, rain, vehicle traffic, pumps, even the tidal effects of the moon and the sun. To overcome these shortcomings, oil explorers tried deploying the tiltmeters deeper in the ground, but that required manual leveling, which could not be done accurately at depths greater than 20 feet (6 meters).

Solving the Problems

Lawrence Livermore researcher David Castillo had a better idea. He and colleagues Carl Boro and Steven Hunter developed the new Oil Field Tiltmeter that is smaller, sends



digital signals, takes less time to install, and costs considerably less than competing products. Working with Pinnacle Technologies in San Francisco, they brought the product to market in October 1996 as the Pinnacle 5000 Oil Field Tiltmeter—and subsequently won an R&D 100 Award for the technology.

"The tiltmeter works on the same principle as a carpenter's level," Hunter says. "The sensor is a liquid-filled glass tube with a gas bubble in it. The difference is that the tiltmeter sensor has electrodes in it so the circuitry can detect the position of the bubble."

The self-leveling mechanisms allowed a smaller tiltmeter design. For deployment at 40- to 100-foot (12- to 30-meter) depths, Hunter and the team designed their product for slim-hole technologies—wells that are only 3 inches (7.6 centimeters) in diameter. In turn, the tiltmeters benefit from more stable hole conditions because the ground is disturbed less when smaller holes are drilled. Other features of the new tiltmeter include an electronic compass for instrument orientation, a downhole analog-digital converter to reduce electronic noise, and an internal data logger with a large memory.

The signal from the new tiltmeter goes to a differential amplifier, is rectified (made to flow in one direction only), and is further amplified. After it is filtered, the signal is digitized by a 24-bit converter, which provides a signal with higher resolution than previously used 16-bit converters. This digitized signal is much less susceptible to electronic noise and thus results in much more reliable data, which are stored in a random-access memory device until a technician downloads it to a lap-top computer.

Pulling It All Together

In an oil field, an array of about 20 instruments defines the magnitude and direction of the tilt vectors at each location; these tilt vectors are used to generate a map of the surface deformation around the well. A modeling program derives the hydrofracture direction that must be present to produce the

(at left) Holding their award-winning Oil Field Tiltmeter are Carl Boro, Steven Hunter, and David Castillo. In a closeup photo of the tiltmeter (below), the green analog card receives electrical signals from the bubble level; below it are the reduction-gear motors that drive the tilting; one of the bubble levels sits on the bottom platform.



observed tilt vectors, and these data enable explorers to determine the best locations for additional wells.

This information is invaluable when an explorer is planning a deep well. The cost of drilling a 10,000-foot-deep well is far greater than twice the cost of drilling a well that is one-half that deep. Such a well will cost at least \$1 million—much more if problems develop during drilling. Given these high development costs, it is essential that an oil exploration company has solid data to rely on.

When compared with other products on the market, an array of the new Pinnacle 5000 instruments shows distinct advantages. Resolution of the angle-of-tilt data has improved from 10 nanoradian to about 1 nanoradian. How fine is this resolution? If you could lay two straight rods next to each other stretching from New York to San Francisco and create an angle by placing a a quarter between them in New York, the angle would be a nanoradian. In fact, this tiltmeter can measure the angle created by adding an atom or two under one side. Power has been reduced from 360 to 200 milliwatts for longer battery life, operating range has been increased from 3 to 10 degrees of tilt for simplified installation, and installation time has been reduced from 40 minutes to 25 minutes. But the most impressive features are the unit's small size, enabling it to fit in a 3-inch-diameter hole vs the 10-inch holes required for other units; the built-in data logger, which allows deeper deployment with less electrical noise; and its low cost.

Hunter says that other applications for the new tiltmeter include tools for studying earthquake faults and geothermal energy production, civil engineering monitors, improved underground mapping for environmental cleanup projects, and more accurate monitoring of subsurface waste disposal operations.

—Sam Hunter

Key words: earthquake faults, geothermal energy, oil-well-drilling, Oil Field Tiltmeter, R&D 100 Award.

For further information contact Steven Hunter
(510) 423-2219 (hunter5@llnl.gov).

Smaller Insulators Handle Higher Voltage



ELECTRONICS advancements throughout the last half century have been rapid, improving lives and commerce in amazing ways. But one technology crucial to many electronics and electrical products—the insulator—has not kept up with these dramatic improvements. Rather, insulator technology has evolved in a steadier, incremental fashion; its improvements have come from using better fabricating materials and reducing manufacturing flaws.

The pace of insulator improvement has now taken a leap forward. Lawrence Livermore researchers led by Stephen Sampayan, together with a research team at AlliedSignal, led by Mike Krogh, have invented the Ultra High Gradient Insulator, or Ultra-HGI, a device that can reliably withstand electrical voltages four times greater than before. That means it will be a smaller, less bulky component in high-tech instrumentation such as accelerators, x-ray machines, semiconductor production tools, and large microwave tubes. Such instrumentation in turn can be designed to be smaller, thereby reducing capital and operating costs. The Ultra-HGI will allow technologies to be advanced in ways never before possible.

How It Works

The Ultra-HGI was developed in response to researchers' needs for smaller insulators; however, smaller size used to mean breakdown at a lower voltage. During the high-voltage conditions that cause a conventional insulator to break down, electrons that are emitted from the insulator surface drift in the electrical field but can impact the surface. This collision liberates a greater number of electrons in a secondary emission, which leaves a net positive charge on the insulator surface that attracts yet more electrons. The result is an escalating cycle of electron collisions and liberations, called a secondary emission avalanche. The avalanche creates a dense plasma that prevents the insulator from sustaining an electric field; the insulator short-circuits and breaks down.

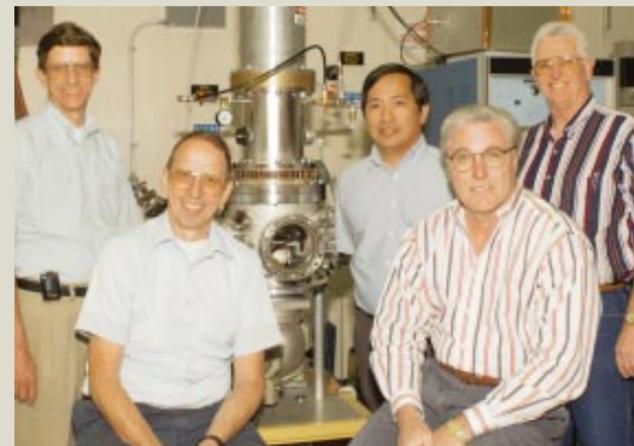
The Livermore insulator designers postulated that an interruption of the electron collisions would avert the avalanche, and breakdowns would not occur. They configured such an interruption by making an insulator with conductive layers that alternate with insulating layers. The

conductive layers are spaced at distances equivalent to electron travel distances, which isolates the insulating layers from electron charges.

The Ultra-HGI is wholly different from today's conventional insulators, which are fabricated as single pieces or as multisection designs. In a single-piece insulator, any flaws make the insulator susceptible to premature failure because the flaw becomes the focus of extremely high-voltage gradients, which initiate and then propagate failure mechanisms. In the conventional multisection insulator, failure in one section causes all the voltage to be redistributed across the few remaining good sections, thereby making them more susceptible to failure. In comparison, the Ultra-HGI's alternating layers divide voltages so finely that the chances of failure are small, and when failures do occur, they are confined to a very small portion of the insulator. An Ultra-HGI insulator comprising a 2.5-millimeter-thick layered stack could have more than 100 conductive "capacitor" layers, each of which receives only a small portion of the high-voltage gradient. In the event of a failure or breakdown of one of the capacitors, the entire voltage is reapportioned across enough remaining capacitors that each receives only a small increase in voltage, and the insulator remains viable. Thus, any breakdown is temporary, no permanent damage is caused, and the insulator can be reconditioned for further use.

Making the Insulator

The prototypes that the research teams made to test their new insulator concept required labor-intensive fabrication. The insulator's layers—which use, for example, copper,



Livermore collaborators surrounding the Ultra High Gradient Insulator are (left to right) Dave Sanders, Ted Wieskamp, Steve Sampayan, Bob Stoddard, and Dave Trimble.

Insulator material	Breakdown threshold, kV/cm	
	Ultra-HGI	Conventional insulator
Polycarbonate	200	50
Fused silica (glass)	175	35
Alumina	125	30

chromium, or aluminum as conductive materials and polycarbonate, glass, or alumina as insulating materials—must be very fine. The thickness of the conducting layers is less than 1 micrometer, while that of the insulating layers is under 1 millimeter. A 1-centimeter thickness of insulator material may contain as many as 40 layers.

One prototype fabrication used perfectly flat, 0.25-millimeter-thick glass plates onto which 0.5-micrometer-thick chromium layers were deposited on both the top and bottom surfaces of the plates. A 2.5- to 3-micrometer layer of gold was added over that.

The metallized plates were aligned and stacked. In a furnace, the stack was subjected to pressure and heated long enough for the plates to form a strong bond. Then the engineers used an ultrasonic abrasive drill to shape the stack into a cylinder and cut a hole through its center. The result was a stack of rings of laminated insulating layers in planes perpendicular to the axis of the cylinder. Each ring was equivalent to a thin high-voltage capacitor; the stack of flat rings constituted a capacitive voltage divider.

The insulator development team tested their prototype by subjecting it to several low-voltage conditioning pulses in a vacuum chamber. They increased the pulses by small amounts until the insulators broke down; the tests were repeated to check the consistency of results. The table above shows the threshold at which the Ultra-HGIs and conventional insulators break down.

Such prototype Ultra-HGIs have been expensive to produce, but the team is now developing designs for mass production and has experimented with several more prototypes of the insulator. While formulating these more producible

designs, they are also experimenting with other ways to prevent or control voltage breakdowns.

Many Present and Future Applications

One reason the insulator was developed was that a new linear accelerator concept proposed by Livermore scientists required a compact insulator. The Dielectric Wall Accelerator will be an order of magnitude smaller than current linear accelerators, but it will deliver similar energy. The narrow separations between the insulator's conductive layers also have been shown to attenuate microwave power, modifying the microwave resonances that cause beam instabilities in accelerators. The Ultra-HGI reduces these resonances by a factor of 4.

While the Ultra-HGI will revolutionize linear accelerators, it will also be important for particle accelerators such as x-ray machines. It should reduce the size—and thus cost—of using such machines for lithography and medicine. It will allow improved performance of high-powered microwaves and neutron sources used for oil-well logging and for detecting explosives. The smaller size and tolerance of higher voltages provided by the Ultra-HGI should make new, smaller designs feasible and economically viable.

—Gloria Wilt

Key Words: capacitor, electron emission, graded insulator, insulator fabrication, R&D 100 Award, secondary emission avalanche, Ultra High Gradient Insulator (Ultra-HGI), voltage breakdown, voltage divider.

For further information contact Bob Stoddard (510) 422-4877 (stoddard1@llnl.gov).

Computer Storage Management Software: The Next Generation

SPEND a little time using personal computers, and you will soon find that faster processing and more memory are the name of the game. "Bigger and faster" are bywords in the supercomputing world as well.

As computers have become more powerful and able to perform ever more complex tasks, data files have rapidly expanded. Unfortunately, data transfer rates and storage system capacities have not kept pace with the expansion of processing and memory capabilities. When preliminary work began at several Department of Energy national laboratories on the Accelerated Strategic Computing Initiative (ASCI), improving high-speed storage became a high priority. ASCI, a major component of DOE's science-based Stockpile Stewardship and Management Program and one of the largest supercomputing projects of all time, will be used to assess the safety, security, and reliability of our nuclear stockpile.

For work in creating an entirely new approach to data storage management and transfer, Lawrence Livermore, Los Alamos, Sandia, and Oak Ridge national laboratories, together with IBM Global Government Industry, recently won an R&D 100 Award. They designed and implemented

the High-Performance Storage System (HPSS), which is capable of managing the storage and data transfer needs of even the most demanding supercomputers and high-speed networks.

The collaboration is based on the premise that no single organization has the ability to confront all of the issues that must be resolved for significant advances in high-performance storage system technology. Lawrence Livermore and its sister DOE laboratories have played leadership roles in this work because of their long history of development and innovation in high-end computing in order to accomplish their national defense and scientific missions.

High-Performance Storage System

The HPSS software is designed to manage data storage capacity into the petabyte range (a quadrillion, or 10^{15} , bytes) and data transfer rates in the gigabyte-per-second range (a billion, or 10^9 , bytes per second). It can move very large data files among high-performance computers, workstation clusters, and storage libraries at speeds 100 to 1,000 times faster than are possible with conventional storage software systems. These speeds support new large-scale applications in high-

performance computing, data collection and analysis, and imaging. For example, high-definition digitized video in real time is now a possibility. The key breakthroughs are a network-centered design, parallel data input/output, and servers that can be distributed and replicated.

All competing products were designed at a time when terabyte (a trillion, or 10^{12} bytes) storage capacities and megabyte- (million-) per-second transfer rates were the target. Unlike most large-scale data storage systems preceding it, HPSS has a network-centered design. In conventional systems, general-purpose computers act as storage servers that connect to storage units such as



Posing among elements of Livermore's new IBM SP computer are the HPSS team members: (front row left to right) Dick Watson, Tammy Dahlgren, and Donna Mecozzi; (back row) Mark Gray, Steve Solbeck, Jim Minton, Norm Samuelson, and Dave Fisher.

disks and tapes. The servers act as intermediaries in passing data to client systems. As data rates increase for storage devices and communications links, the size of the server must also increase to provide the required capacity and total data throughput bandwidth. These high data rates and capacity demands tend to drive the storage server into the mainframe class, which can be expensive to purchase and maintain.

If the storage software system and storage devices are instead distributed over a network, control of the storage system can be separated from the flow of data. The bottleneck is removed, allowing more rapid data transmission and expanded performance and capacity. Workstation-class systems used as storage servers provide the high performance required and reduce the cost for storage-server hardware in the bargain.

Focus on the Network

Operating on a high-performance network, the HPSS is designed to allow data to be transferred directly from one or more disk or tape controllers to a client. The HPSS accommodates the simultaneous transfer of parallel data streams from multiple storage devices to computers exercising parallel applications. As the speeds and capacities of storage media and devices increase, this software will easily assimilate them to produce transfer rates in the range of multiple gigabytes per second. For example, if a system has a storage device that can deliver 100 megabytes per second but a gigabyte per second is needed, then 10 devices in parallel, controlled by HPSS software, can be used to expand, or "scale up," to the new requirement. With this design, the HPSS will be able to handle almost unlimited storage capacity, data transfer rates of billions of bytes per second and beyond, virtually unlimited file sizes, millions of naming directories, and hundreds to thousands of simultaneous clients.

HPSS uses several mechanisms, including "transactions," to ensure data reliability and integrity. Transactions are groups of operations that either take place together or not at all. The problem with distributed servers working together on a common job is that one server may fail or not be able to do its part. Transactions assure that all servers successfully complete their job or the function is aborted. Transactional integrity is common in relational data management systems, but it is new in storage systems.

HPSS is designed to support a range of supercomputing and multiprocessor client platforms, operate on many vendors' platforms, and use industry-standard storage hardware. The basic infrastructure of HPSS is the Open Software Foundation's Distributed Computing Environment because of its wide adoption among vendors and its almost universal acceptance by the computer industry.

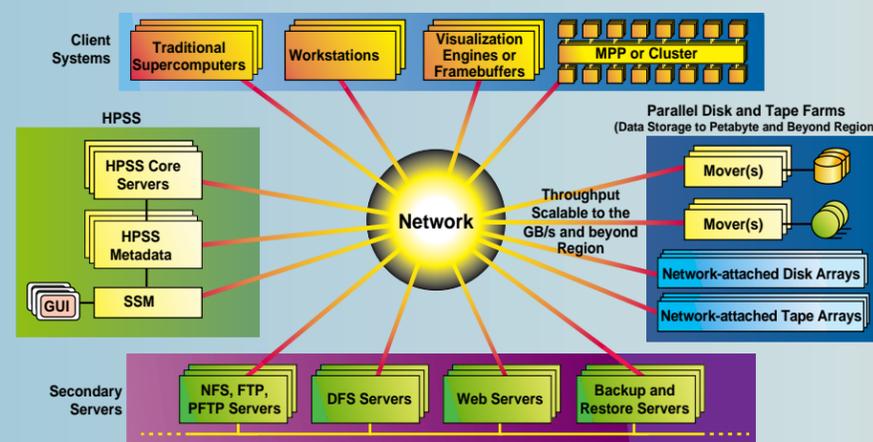
IBM began marketing the system commercially in the fall of 1996. The HPSS has already been adopted by the California Institute of Technology/Jet Propulsion Laboratory, Cornell Theory Center, Fermi National Accelerator Laboratory, Maui High-Performance Computer Center, NASA Langley Research Center, San Diego Supercomputer Center, and the University of Washington, as well as by the participating Department of Energy laboratories. Lawrence Berkeley National Laboratory has also adopted the system through its National Energy Research Supercomputer Center, which was an original participant in the HPSS work.

There are a number of other prospective users of the HPSS. A possible customer is considering digitizing its entire film archive to produce several petabytes of data. Other applications might include oil company databases, the finance and insurance sector, the Human Genome Project, medical imaging and records, and high-energy physics. In combination with computers that can produce and manipulate huge amounts of data at ever-increasing rates, the HPSS expandable, parallel, network-based design gives users the capability to solve problems and manage information more easily than ever.

—Katie Walter

Key Words: Accelerated Strategic Computing Initiative (ASCI), computer network, hierarchical storage management, High-Performance Storage System (HPSS), large-scale computer storage, parallel computing, R&D 100 Award, supercomputing.

For further information, contact Dick Watson (510) 422-9216 (dwatson@llnl.gov), or visit the HPSS Internet home page at <http://www.sdsc.edu/hpss/>.



The High-Performance Storage System connects the storage units directly to a high-speed network. This network-centered architecture allows the data to flow directly between client systems and storage units, bypassing the storage server. Throughput may exceed the rate of a billion bytes per second.

Each month in this space we report on the patents issued to and/or the awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory.

Patents

Patent issued to	Patent title, number, and date of issue	Summary of disclosure
Steven T. Mayer Richard W. Pekala James L. Kaschmitter	Composite Carbon Foam Electrode U.S. Patent 5,626,977 May 6, 1997	Composite carbon foam electrodes that incorporate granularized materials. The composite carbon foam is produced by pyrolysis of mixtures of granularized material, polyhydroxybenzene compounds, and formaldehyde or furfural. The foam has a surface area of 100 to 1200 m ² /g and a density of 0.1 to 1.2 g/cm ³ . The granularized materials include metal and/or carbon fibers, carbon aerogel microspheres, crushed or powderized carbon aerogel, powderized carbon, and metal and metal oxide powders and spheres. The granularized materials are added during the phase when the precursor materials are liquid (prior to gelation), and the precursor materials are spread in very thin films or as thin films on roll-to-roll substrates.
Thomas E. McEwan	Micropower RF Transponder with Superregenerative Receiver and RF Receiver with Sampling Mixer U.S. Patent 5,630,216 May 13, 1997	A radio-frequency (rf) receiver that includes an external quench oscillator for generating a series of pulses at a predetermined quench frequency and a pulse-forming network connected to the quench oscillator for converting the series of periodic pulses into a series of exponentially damped drive pulses. An oscillator is connected to and driven by these drive pulses, and an antenna is connected to the oscillator for receiving modulated rf signals. A signal extraction network is also connected to the oscillator for blocking the quench frequency signals and for passing the detected rf signals. A micropower amplifier is connected to the signal extraction network for amplifying the detected rf signals. Alternatively, an rf receiver includes a receive antenna, a frequency-selecting network connected to the antenna, and a sampling mixer connected to the network.
Thomas J. Karr	Ballistic Projectile Trajectory Determining System U.S. Patent 5,631,654 May 20, 1997	A system that determines the three-dimensional trajectory of a ballistic projectile. To initialize the system, predictions of state parameters for a ballistic projectile are received at an estimator. The estimator uses the predictions to estimate first trajectory characteristics of the ballistic projectile. A single stationary monocular sensor then observes the actual first trajectory characteristics of the ballistic projectile. A comparator generates an error value related to the predicted state parameters by comparing the estimated first trajectory characteristics of the ballistic projectile with the observed first trajectory characteristics of the ballistic projectile. If the error value is equal to or greater than a selected limit, the predictions of the state parameters are adjusted. The process is repeated until the error value is less than the selected limit. A computer then calculates trajectory characteristics.
Leland B. Evans Vincent Malba	Pressure Activated Diaphragm Bonder U.S. Patent 5,632,434 May 27, 1997	An apparatus constructed to bond integrated circuit chips to a substrate to form multichip modules or hybrid packages, using thin-film eutectic solders or bonding materials. The bonder apparatus uses a diaphragm that allows a uniform, high pressure to be applied to the chips, without causing chipping or cracking, while fracturing any solder oxide and providing solder flow. The bonder also uses a template that provides highly accurate chip placement on the substrate and can simultaneously bond chips of different sizes and thicknesses. The bonder also can operate in the presence of any gas (nitrogen, helium, forming gas) or under vacuum, thus enabling a controlled environmental operation.
Alexander R. Mitchell Philip F. Pagoria Robert D. Schmidt	Vicarious Nucleophilic Substitution Using 4-Amino-1,2,4-Triazole, Hydroxylamine or O-Alkylhydroxylamine to Prepare 1,3-Diamino-2,4,6-Trinitrobenzene or 1,3,5-Triamino-2,4,6-Trinitrobenzene U.S. Patent 5,633,406 May 27, 1997	A process to inexpensively produce the insensitive benzene ring-based explosives 1,3-diamino-2,4,6-trinitrobenzene (DATB) and 1,3,5-triamino-2,4,6-trinitrobenzene (TATB). 4-amino-1,2,4-triazole (ATA), as well as hydroxylamine or its O-alkyl derivatives, are used to synthesize DATB and TATB by vicarious nucleophilic substitution reactions. A trinitroaromatic starting material is reacted with ATA in the presence of a strong base in a solvent to produce DATB or TATB. Starting materials include 1,3,5-trinitrobenzene, 2,4,6-trinitroaniline, or 1,3-diamino-2,4,6-trinitrobenzene and may be produced by various reactions. Strong bases include sodium methoxide or potassium tert-butoxide, and a suitable solvent is dimethylsulphoxide (DMSO). The reaction temperature is 10 to 30°C. Yields of over 90% can be obtained.

Patents (continued)

Patent issued to	Patent title, number, and date of issue	Summary of disclosure
James L. Kaschmitter Tri D. Tran John H. Feikert Steven T. Mayer	Fabricating Solid Carbon Porous Electrodes from Powders U.S. Patent 5,636,437 June 10, 1997	The fabrication of conductive solid porous carbon electrodes for use in batteries, double-layer capacitors, fuel cells, capacitive deionization, and waste treatment. Carbon-containing powder, such as pyrolyzed carbon aerogel microspheres or powder, and a binder precursor, such as a phenolic resin, are formed into a slurry to which a dopant may also be added. A film is formed by depositing the slurry on a substrate or on a porous material. The film is pyrolyzed in an inert, reducing, or oxidizing atmosphere at 600°C to form an electrode. The composition of the electrode determines its use.
David J. Erskine	White Light Velocity Interferometer U.S. Patent 5,642,194 June 24, 1997	A technique that allows the use of broadband and incoherent or "white" illumination to measure the velocity of objects. The principle of white light velocimetry can be applied to any wave phenomenon, i.e., light as well as microwaves and sound. Powerful, compact, or inexpensive sources can be used for remote target velocimetry. These include flash and arc lamps, light from detonations, pulsed lasers, chirped-frequency lasers, and lasers operating simultaneously in several wavelengths. A double interferometer system uses two broadband (achromatic), superimposing interferometers in series with an interposed target. Light rays of one interferometer arm are delayed in time but superimposed in path for a wide range of wavelengths.

© 1997. The Regents of the University of California. All rights reserved. This document has been authored by the The Regents of the University of California under Contract No. W-7405-Eng-48 with the U.S. Government. To request permission to use any material contained in this document, please submit your request in writing to the Technical Information Department, Publication Services Group, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or to our electronic mail address report-orders@llnl.gov.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California and shall not be used for advertising or product endorsement purposes.

U.S. Government Printing Office: 1997/583-059-60027